Network Working Group
INTERNET DRAFT
Expires: 10 July 1994
<draft-ietf-atm-framework-doc-00.ps>

R. G. Cole
AT&T Bell Laboratories
11 January 1994

IP over ATM: A Framework Document

Status of this Memo

1. Abstract

This document summarizes some of the discussions of the *IP over ATM* working group over the last year. This summary is provided in the form of a framework which categorizes the various ATM subnet models and End-to-End models identified by the working group. The intent of this framework is to help partition the efforts of the working group to focus on smaller, more manageable aspects of *IP over ATM*. This is deemed necessary due to the number of ATM subnet types being discussed.

The subnet models identified in this memo are:

- an SVC-based LAN ATM (LATM) subnet,
- a PVC-based WAN ATM (WATM) subnet, and
- an SVC-based WATM subnet.

The End-to-End models identified in this memo are:

- End-to-End Subnet Models, including
  - — the Classical IP model,
  - — an SVC-based WATM subnet model, and
  - — lightweight subnet models, i.e. TUNIC and TULIP, and
- Peer Models.

2. Introduction

The *IP over ATM* Working Group of the Internet Engineering Task Force (IETF) is chartered to develop standards for routing and forwarding IP packets over ATM subnets. There are several difficulties with this charter, as witnessed in the discussions of this group over the last several years. First, asynchronous transfer mode (ATM) technology is still in its development state. Second, the flexibility of the technology, allows and in fact encourages, viewing ATM as different subnet technologies, e.g., a LAN, a WAN, an Internet, etc. Therefore, much of the discussions of the *IP over ATM* WG have centered around not only how to carry IP traffic over ATM but what is ATM. This memorandum attempts to highlight some of these discussions and tries to categorize the various ATM subnet architectures and End-to-End architectures that have been proposed.

It is felt that to promote progress in the *IP over ATM* working group, it is first necessary to come to some consensus on the types of ATM subnets to be addressed. Then it will be possible to focus attention of several subgroups to work the issues of routing IP packets over these different ATM subnets. As this work is proceeding, parallel efforts in identifying issues to work in the eventual end-to-end ATM models can be performed. It is expected that defining IP routing over some ATM subnet types and simpler end-to-end models can proceed at a quicker pace, while there exist many fundamental architectural issues in defining peer end-to-end models. Therefore, it is hoped that this document, in classifying the subnet and end-to-end models, will help in focusing the *IP over ATM* working group's direction over the next several IETF meetings.

The remainder of this memorandum is organized as follows:

- Section4 presents several definitions of use in the later sections,
- Section 5 discusses/identifies the various subnet models proposed to date by the *IP over ATM* working group,
- Section 6 identifies several promising end-to-end networking configurations, and
- Section 7 contains conclusions.

3. PARAMETERS DEFINING SUBNET TECHNOLOGIES

Following ISO we define several terms **[ISO 8473]**:

A *Subnet* is a connected communication network consisting of a single networking technology. Subnets are further categorized into *broadcast* and *general topology* subnetworks.

A *Broadcast Subnet* "supports an arbitrary number of end systems and intermediate systems and additionally is capable of transmitting a single Subnetwork NPDU (SNPDU) to a subset of these systems in response to a single send_unitdata request".

A *General Topology Subnet* "supports an arbitrary number of end systems and intermediate systems but does not support a convenient multi-destination connectionless transmission facility, as does a broadcast subnetwork". We will hereafter interchangeably use the terms general topology subnet and Non-Broadcast Multiple Access (NBMA) subnet.

An *End System* "delivers Network Level Protocol Data Units (NPDUs) to other systems and receives NPDUs from other systems, but do not relay NPDUs."

An *Intermediate System* "delivers and receives NPDUs from other systems, and relays NPDUs from other systems to other destination systems." They route to end systems within their own area and to other intermediate systems when the destination is within another area.

An *End-to-End path* consists of two end systems which can communicate with one another over an arbitrary number of intermediate systems and subnets connecting the end and intermediate systems.

From the perspective of the subnet, it sees no difference between the end systems and intermediate systems. Therefore, we will define a Subnet End Systems (Sub-ES) as a network level routing entity connected to the subnet, whether it is actually an end or intermediate system. Figure 1 shows two subnet end systems connected over a single subnet. An IP-level (or L3 level) end-to-end path can be considered as a concatenation of

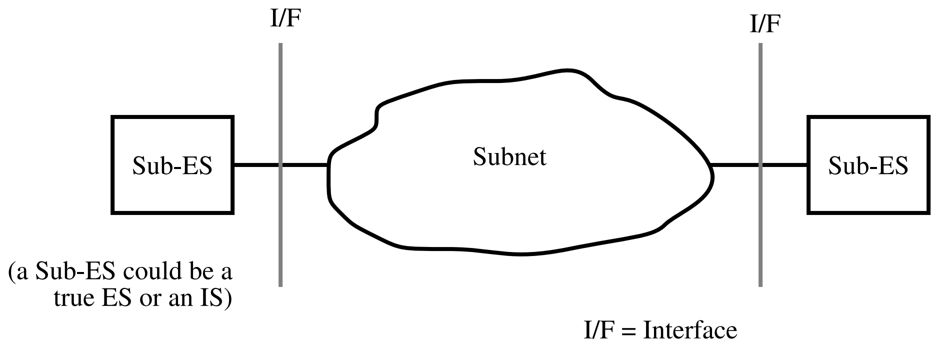

Figure 1: A Single Hop configuration between Subnet End Systems (Sub-ES) attached to the same subnet

multiple IP-level (or L3) single hop paths.

Each subnet technology is characterized in terms of several parameters. We have defined two subnet types, i.e. a broadcast and a general topology, or NBMA, subnet. However, it is possible to further specify subnet types within these broad categories. Towards this end, we list here a few of the more relevant parameters for the discussions to follow; this list is by no means exhaustive.

They are:

- Routing and Addressing Capabilities,
- Topology and Geographic Support, e.g.,
  - — LAN versus WAN, and
  - — the number of stations supported,
- Transport Services, e.g.,
  - — Connection-Oriented Network Service (CONS) versus Connectionless Network Service (CLNS), and
  - — Variable Bit Rate (VBR) versus Constant Bit Rate (CBR) Services,
- Connection Functions, e.g.,
  - — Permanent Virtual Connection (PVC) or Switched Virtual Connection (SVC), and
  - — point-to-point, multicast, or broadcast,
- Transmission Characteristics, e.g.,
  - — reliable versus unreliable (e.g., guaranteed delivery as in X.25 or best effort as in Frame Relay),
  - — priorities, and

— qualities of service,

- Security Features, and
- Packet Parameters (e.g., max MTU size, etc.).

4. CLASSIFICATION OF ATM SUBNETS

Several ATM subnets have been discussed by the *IP over ATM* working group. These range from the Broadband-Integrated Services Digital Network (B-ISDN) subnet originally proposed and being defined by interexchange carriers to Local ATM (LATM) subnets, currently being defined by computer and LAN networking vendors. Further, other subnets are defined on ATM as overlay subnets, e.g. Frame Relay on ATM.

We identify and discuss here three possible ATM subnets; a LAN-ATM (LATM) subnet, a PVC-based WAN-ATM (WATM) subnet, and a SVC-based WATM subnet.

4.1 The LATM Subnet Model

The LAN-ATM (LATM) subnet is characterized as having a small number of end systems, e.g., up to at most a few thousand, which:

- are directly connected,
- share a common IP-level (or L3) network prefix,
- support switched virtual connections (SVCs) as well as permanent virtual connections (PVCs),
- support an efficient broadcast/multicast/connectionless server for address resolution, and
- are part of a single working group, e.g. security is not an issue. (Note: An assumption for the LATM is that all end systems attached to the LATM are part of the same administrative domain and that the ATM network will setup up connections between any pair of end systems. Connectivity to the “outside” world is achieved only through a router which will provide the firewall if so desired.)

For end systems to forward packets over this subnet, many of the mechanisms for forwarding IP packets over ethernet subnets apply and are presented in **[LAU 93]**. Issues to be worked in defining IP over the LATM subnet are:

- Defining the VPI/VCIs for address resolution, (Many internal architectures are possible ranging from, possibly, a broadcast capability to the existence of an ARP server reachable on the identified VPI/VCI. It is the belief of the *IP over ATM* group that the architecture of address resolution needs to be specified and that the ARP server architecture is the desired option. The APR server is specified in **[LAU 93]**.)
- Defining the maximum MTU size, (This has been the subject of considerable debate over the Internet this spring. This issue is addressed in **[ATK 93]**.)
- Methods of encapsulation and multiplexing, (This issue is addressed in **[HEI 93a]**.)
- Reliability of the ARP server mechanism defined in **[LAU 93]**,
- Support for IP multicasting,
- Defining SVC optimization techniques, e.g., timeout mechanisms,
- Negotiations of, e.g., MTU size, method of encapsulation, and

- Special purpose SVC for control/routing and high bandwidth data transfers.

The majority of the work here is in defining the interface between the Sub-ESs and the LATM, e.g., what VPI/VCIs to be used to access the ARP server **[LAU 93]**, the method of encapsulation, the maximum MTU sizes, signalling capabilities, etc. However, other work items involve more internal architectural issues such as the architectures for ARP services.

4.2 The PVC WATM Subnet Model

The PVC-based WATM is a general topology, or NBMA, subnet which is further characterized as having a large number of end systems, e.g., up to at most a few tens of thousands, which:

- are not all directly connected,
- do not share a common IP-level (or L3) network prefix,
- support permanent virtual connections (PVCs),
- do not have to support an efficient broadcast/multicast/connectionless server for address resolution, and
- are not part of a single working group, however security is not an issue due to the permanent nature of the virtual connections.

For end systems to forward packets over this subnet, many of the mechanisms defined for forwarding IP packets over Frame Relay subnets apply, see **[RFC 1490]**. Issues to be worked in defining IP over the LATM subnet are:

- Defining the details for an Inv-ARP like address resolution, (Strictly, Inv-ARP is defined for general PVC-based NBMA subnets, see **[RFC 1293]** and **[LAU 93]**.)
- Methods of encapsulation and multiplexing, (This issue is addressed in **[HEI 93a]**.)
- Defining the maximum MTU size, (This issue is addressed in **[ATK 93]**.) and
- Support for IP multicasting.

The majority of the work here is in specifying the Inv-ARP details and in defining the MTU sizes supported.

4.3 The SVC WATM Subnet Model

The SVC-based WATM is a general topology, or NBMA, subnet which is further characterized as having a large number of end systems, e.g., up to at most a few hundreds of thousands, which:

- Support switched virtual connections (SVCs) as well as permanent virtual connections (PVCs),
- Maybe directly connected, due to the existence of SVCs,
- Do not share a common IP-level (or L3) network prefix,
- Support for multiple subnetwork point of attachment (SNPA) address formats (four possible formats are allowed in **[ATM 1993]**),
- Support negotiations for, e.g., MTU size, methods of encapsulation,
- Do not have to support an efficient broadcast/multicast mechanism for address resolution,
- Are not part of a single working group and hence security is an issue, and
- Are potentially billed for holding circuit holding time.

For end systems to forward packets over this subnet, many of the problems which were being addressed in the IETF's *IPLPDN* working group, the *IP over ATM* working group, and are now to be addressed in the *Routing over Large Clouds* **[HAL 93]**, are applicable. Examples of work performed in the *IPLPDN* working group include short-cut routing **[TSU 93]** and directed ARP **[GAR 93]** over SMDS subnets, and in the *IP over ATM* working group include distributed ARP server architectures in **[HEI 93b]**. Issues to be worked for the SVC WATM subnet include:

- Defining the details for an efficient address resolution architecture, (Work in this area has been started, see **[GAR 93]**, and **[HEI 93b]**.)
- Methods of encapsulation and multiplexing, (This issue is addressed in **[HEI 93a]**.)
- Defining the default MTU size, (This issue is addressed in **[ATK 93]**.)
- Defining security procedures,
- Routing between end systems not sharing a common IP-level (or L3) network prefix, but attached to the same subnet,
- Support for IP multicasting,
- Defining SVC optimization techniques, e.g., timeout mechanisms,
- Negotiations of, e.g., MTU size, method of encapsulation, and
- Special purpose SVC for control/routing and high bandwidth data transfers.

Whereas the first two subnets presented above, i.e. the LATM and PVC-Based WATM, are more straight forward in defining routing of IP packets over these subnets, the SVC-Based WATM presents several hard problems to be resolved prior to specifying IP routing over it. Several of these issues are the focus of the new *Routing over Large Clouds* working group formed in July 1993.

5. END-TO-END NETWORKING CONFIGURATIONS

Several end-to-end ATM configurations have been proposed; most constructed by cascading gateways and different ATM subnet models. Other, more radical models have been proposed which cast out the traditional subnet architectures in favor of ATM Internet models. We categorize these different models in terms of whether the ATM network and IP-level (or L3) routers are considered as peers, i.e. they exchange routing information, or not, i.e. the IP-level (or L3) routers treat the ATM network as a lower-level subnet. We will refer to these as *Peer Models* and *Subnet Models*, respectively, and discuss each separately below.

Clearly the Subnet Models follow the traditional IP networking architectures, where as the Peer Models propose a new way of IP networking. Due to this, the consensus of the group was that the problems with the end-to-end peer models were much harder than any of the other models, and had more unresolved technical issues. While encouraging interested individuals/companies to research this area, it was not the priority of the *IP over ATM* working group to address these difficulties.

5.1 Subnet Models

Three end-to-end subnet models were identified at the Columbus (March/Spring 1993) and Amsterdam (July/Summer) IETF meetings of the *IP over ATM* working group. These three models are

- The Classical IP Model,
- The SVC-based WATM subnet model, and
- Lightweight subnet models.

5.1.1 The Classical IP Model

The Classical IP Model was suggested at the Spring 1993 IETF meeting **[LAU 93]**. This model simply consists of cascading subnet models discussed above with IP-level (or L3) routers. A example realization of this model consists of a concatenation of two (simplified) LATMs and a PVC-based WATM subnet. This is shown in Figure 2 below. Once the details of specifying routing and forwarding IP packets over the component subnets are complete, routing IP packets over this Classical IP model is straight forward. The LATMs are simplified in that they:

- support a single subnetwork point of attachment (SNPA) address format (four allowed address formats are specified in **[ATM 1993]**,
- support a single default MTU size, and
- support a default LLC/SNAP encapsulation as specified in **[HEI 93a]**.

For more information see **[LAU 93]**.

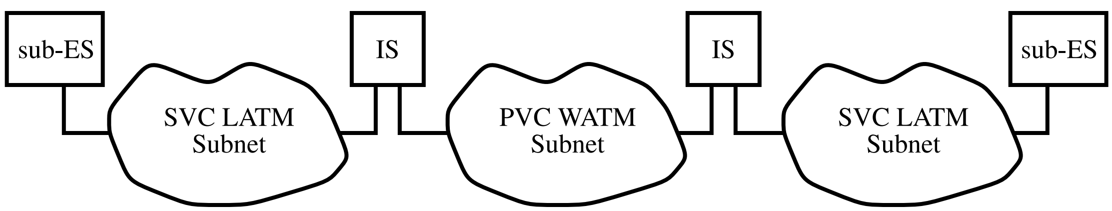

Figure 2: The Classical IP model as a concatenation of three separate ATM subnets; two SVC LATM subnets and a single PVC WATM subnet.

5.1.2 The SVC-based WATM subnet

The SVC-based WATM subnet was discussed in Section 5.3. Here the subnet supports a huge number of ESs in a dynamic SVC environment. There are several problems to address to bring this end-to-end subnet model to fruition. These include, but are not limited to, the following:

- What are the appropriate architectures for address resolution, and
- How can routing be optimized across multiple logical IP subnets on a common ATM infrastructure, e.g. the short-cut routing problem.

The work on these problems may be worked separately from the near term work on bringing the Classical IP model to fruition. Figure 3 shows an example of an SVC-based WATM subnet consisting of three components, two SVC LATM networks connected by a single PVC or SVC WATM network. More complex ATM internets are envisioned and solutions to networking over SVC-based WATM subnets must be able to support these complex internets and their attending problems. Many of the routing and other internetworking issues associated with this model are now to be addressed in the newly formed *Routing over Large Clouds* working group.

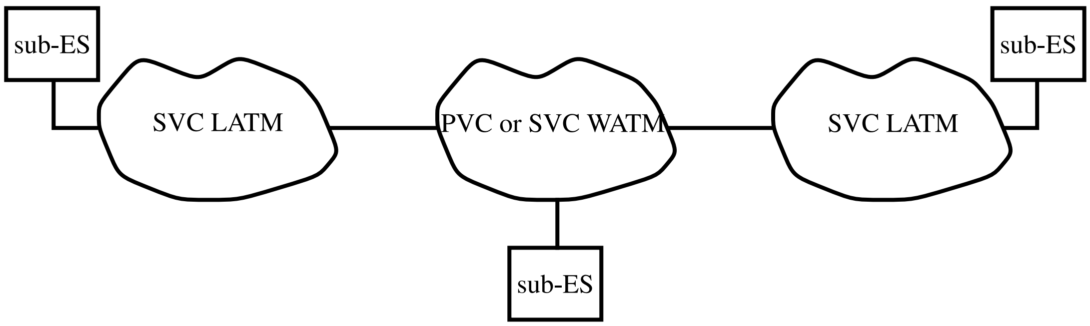

Figure 3: An SVC-based WATM configuration between End Systems which in fact could be ESs or ISs.

An additional complexity in supporting IP routing over these ATM internets lies in the multiplicity of subnetwork points of attachment address formats in **[ATM 93]**. NSAP modeled address formats only are supported on "private ATM" networks, while either 1) E.164 only, 2) NSAP modeled formats only, or 3) both are supported on "public ATM" networks. Further, while both the E.164 and NSAP modeled address formats are to be considered as subnetwork points of attachment, it seems that E.164 only networks are to be considered as subordinate to "private networks", in some sense. This leads to some confusion in defining an ARP mechanism in supporting all combinations of end-to-end scenarios (refer to the discussion in Appendix A on the possible scenarios to be supported by ARP).

5.1.3 Lightweight subnet models

Two further models can be identified which have the characteristic of largely or totally eliminating IP header overheads. These models were discussed in the July IETF meeting in Amsterdam. Given a single hop configuration between Sub-ES as shown in Figure 1, they both assume that following name resolution, address resolution, and SVC signaling, an implicit binding is established between entities in the two subnet ESs. In this case full IP headers (and in particular source and destination addresses) are not required in each data packet.

- The first model is "TCP and UDP over Lightweight IP" (TULIP) in which only the IP protocol field is carried in each packet, everything else being bound at call set-up time. In this case the implicit binding is between the IP entities in each subnet ES. Since there is no further routing problem once the binding is established, since AAL5 can indicate packet size, since fragmentation cannot occur, and since ATM signaling will handle exception conditions, the absence of all other IP header fields and of ICMP should not be an issue.

  TULIP changes nothing in the abstract architecture of the IP model, since each subnet ES still has an IP address which is resolved to an ATM address. It simply uses the point-to-point property of VCs to allow the elimination of some per-packet overhead. The use of TULIP could in principle be negotiated on a per-SVC basis or configured on a per-PVC basis.

- The second model is "TCP and UDP over a Nonexistent IP Connection" (TUNIC). In this case no network-layer information is carried in each packet, everything being bound at virtual circuit set-up time. The implicit binding is between two applications using either TCP or UDP directly over AAL5 on a dedicated VC. If this can be achieved, the IP protocol field has no useful dynamic function. However, in order to achieve binding between two applications, the use of a well-known port number in classical IP or in TULIP mode may be necessary during call set-up. This is a subject for further study.

5.2 Peer Models

Peer Models are fundamentally characterized in that they place routers/gateways and the ATM network on a peer basis. That is, the ATM network and the attached ESs or ISs exchange routing information on a peer basis, network level addressing is used across the ES or IS and ATM interface. Discussions on peer models can be found in **[LYO 92]**, **[HEI 1992]** and **[LIA 93]**.

A rationale for this model is that due to the envisioned complexity of future ATM networks, the routing complexities will be similar to routing over complex internets. Rather than building in redundancy in routing complexities at the network and subnetwork levels, perhaps it is possible to imbed within the ATM network fabric IS functionality. This is illustrated in Figure 4 below. Here the term “Single ATM Domain” refers to a collec-

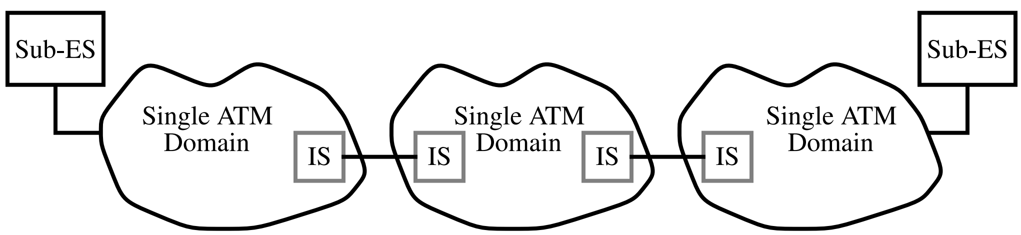

Figure 4: The ATM Internet model imbeds within the ATM cloud network level routing features which in essence peers with traditional gateways or routers.

tion of ATM switches from a single ATM switch vendor or a collection of ATM switches administered by a single authority. The routing function of the IS is only required during the connection set-up phase and is not in the path of the data flow following the connection establishment.

There are issues to be worked for this Peer Model (see, e.g., **[HEI 1992]** and **[LIA 93]**), including:

- Defining a routing architecture which
  - — scales to the large size expected of the eventual ATM network, and
  - — supports TOS routing, and
- Defining a functional architecture which
  - — scales to the large size expected of the eventual ATM network,
  - — supports the current status of the ATM Forum’s standards in signalling, routing, etc., and
  - — allows a transition from subnet models to peer models.

During the discussions of the *IP over ATM* working group, it was felt that the problems with the end-to-end peer model were much harder than any other model, and had more unresolved technical issues. While encouraging interested individuals/companies to research this area, it was not a priority of the working group to address these issues.

5.3 Transition Models

Finally, it is useful to consider transition models, lying somewhere between the Classical IP Models and the Peer Models. Some possible architectures for transition models are presented in **[LIA 93]**. Others are possible, for example Figure 5 showing a Classical IP transition model which assumes the presence of gateways between ATM subnets and ATM Peer networks. Other transition options are discussed in **[LIA 93]**.

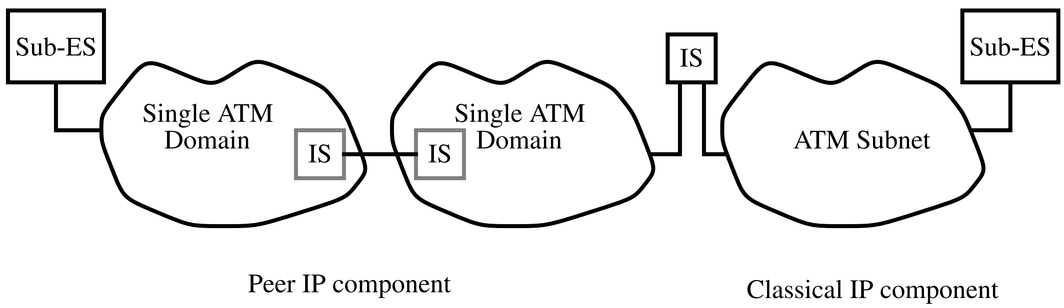

Figure 5: The ATM transition model assuming the presence of gateways or routers between the ATM subnets and the ATM peer networks.

6. Range of Application of the Working Group's Documents

The *IP Over ATM* Working Group has generated several Internet-Drafts or RFCs. This section identifies the relationship of these documents to the various *IP Over ATM* Models identified in this document. The Draft and RFCs produced to date are the following references, **[HEI 1993b]**, **[LAU 1993]**, and **[ATK 1993]**. The following table cross references these application of these document to the various *IP Over ATM* Models. An "x" in the column under a particular reference indicates that the specifications in the reference applies to the particular *IP Over ATM* Model.

Table 1: Cross Reference of WG Documents and IP Over ATM Models

Models/References	[HEI 1993b]	[LAU 1993]	[ATK 1993]	
Subnet Models:				
SVC-based LATM	X	X	X	
PVC-based WATM	X	X	X	
SVC-based WATM	X		X	
End-to-End Models:				
Classical IP	X	X	X	
SVC-based WATM	X		X	
TULIP	X		X	
TUNIC				
Peer	X		X	

7. CONCLUSIONS

This document identifies several of the possible ATM subnet technologies that have been bantered about at the meetings of the *IP over ATM* working group at the IETF. In particular, an LAN-ATM (LATM), a PVC-based WAN ATM (PVC-based WATM), and an SVC-based WAN ATM (SVC-based WATM) are presented. Further, this document lists several of the end-to-end *IP over ATM* architectures which have been discussed and the distinguishing characteristics of each. End-to-end models discussed include SVC-based WATM, light weight IP models, e.g., TULIP and TUNIC models, and Peer models.

It is proposed that the issues of routing *IP over ATM* be broken out into issues relating separately to the subnet and end-to-end models discussed in this document. Due to the complexity of ATM technologies, in that they present multiple faces, it is felt that the best way to make progress is to simplify the work into more manageable steps.

8. Acknowledgments:

This draft is the direct result of the numerous discussions of the *IP over ATM* Working Group of the Internet Engineering Task Force. The author also had the benefit of several private discussions with H. Nguyen of AT&T Bell Laboratories. Brian Carpenter of CERN was kind enough to contribute the TULIP and TUNIC sections to this draft. The text of Appendix A was pirated liberally from Anthony Alles' of HLS posting on the *IP over ATM* discussion list (and modified at the author's discretion). This draft also has benefitted from numerous suggestions from John T. Amenyo of ANS.

9. Author's Address:

Robert G. Cole
AT&T Bell Laboratories
101 Crawfords Corner Road, Rm. 1F-408
Holmdel, NJ 07733
Phone: (908) 949-1950
Fax: (908) 949-1726
Email: rgc@qsun.att.com

10. References:

[ATM 1993] ATM Forum, "ATM User-to-Network Interface (UNI) Specification: Version 3.0", Sept. 1993.

[ATK 1993] R. Atkinson, "Default IP MTU for use over ATM Adaptation Layer 5 (AAL 5)", IETF Network Working Group INTERNET DRAFT, 16 November 1993.

[GAR 1993] J. Garrett, "Directed ARP", IETF Network Working Group RFC 1433, March 1993.

[HAL 1993] J. Halpern, postings to the *IP over ATM* and *Routing over Large Clouds* discussion lists.

[HEI 1993a] J. Heinanen and R. Govindan, "Next Hop Resolution Protocol (NHRP)",*IETF Routing over Large Clouds* Working Group INTERNET DRAFT, 15 October 1993.

[HEI 1993b] J. Heinanen, "Multiprotocol over Adaptation Type 5 on ATM", IETF Network Working Group RFC 1483, July 1993.

[HEI 1992] J. Heinanen, "Routing and Addressing in ATM Networks", private memorandum posted to the *P over ATM* Working Group, 27 November, 1992.

[ISO 8473] "Information technology - Telecommunications and Information exchange between systems - Intermediate system to Intermediate system Intra-Domain routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)," International Organization for Standardization, 1992.

[LAU 1993] M. Laubach, "Classical IP and ARP over ATM", IETF Network Working Group INTERNET DRAFT, 14 October 1993.

[LIA 1993] F. Liaw, "*IP over ATM*: architecture, address translation, and call control", *IP over ATM* Working Group INTERNET DRAFT, 21 March, 1993.

[LYO 1992] T. Lyon, F. Liaw, and A. Romanow, "Network Layer Architecture for ATM Networks", private memorandum, 1992.

[RFC 1293] T. Bradley and C. Brown, "Inverse Address Resolution Protocol", IETF Networking Working Group RFC 1293, January 1992.

[RFC 1293] T. Bradley, C. Brown, and A. Malis, "Multiprotocol Interconnect over Frame Relay", IETF Networking Working Group RFC 1490, July 1993.

[TSU 1993] P. Tsuchiya, "Shortcut Routing: Discovery and Routing over Large Public Data Networks", IETF Network Working Group INTERNET- DRAFT, January 1993.

11. Appendix A: Potential Interworking Scenarios to be Supported by ARP

The architectural model of the VC routing protocol, being defined by the Private Network-to-Network Interface (P-NNI) working group of the ATM Forum, categorizes ATM networks into two types:

- Those that participate in the VC routing protocols and use NSAP modeled addresses **[ATM 93]** (referred to as private networks, for short), and
- Those that do not participate in the VC routing protocol. Typically, but possibly not in all cases, public ATM networks that use native mode E.164 addresses **[ATM 93]** will fall into this later category.

The issue for ARP, then is to know what information must be returned to allow such connectivity. Consider the following scenarios:

- Private host to Private Host, no intervening public transit network(s): Clearly requires that ARP return only the NSAP modeled address format of the end host.
- Private host to Private host, through intervening public networks: In this case, the connection setup from host A to host B must transit the public network(s). This requires that at each ingress point to the public network that a routing decision be made as to which is the correct egress point from that public network to the next hop private ATM switch, and that the native E.164 address of that egress point be found (finding this is a VC routing problem, probably requiring configuration of the public network links and connectivity information). ARP should return, at least, the NSAP address of the endpoint in which case the mapping of the NSAP addresses to the E.164 address, as specified in **[ATM 93]**, is the responsibility of ingress switch to the public network.
- Private Network Host to Public Network Host: To get connectivity between the public node and the private nodes requires the same kind of routing information discussed above - namely, the directly attached public network needs to know the (NSAP format) ATM address of the private station, and the native E.164 address of the egress point from the public network to that private network (or to that of an intervening transit private network etc.). There is some argument, that the ARP mechanism could return this egress point native E.164 address, but this may be considered inconsistent for ARP to return what to some is clearly routing information, and to others is required signaling information.

  In the opposite direction, the private network node can use, and should only get, the E.164 address of the directly attached public node. What format should this information be carried in? This question is clearly answered, by Note 9 of Annex A of **[ATM 93]**, Version 2.4 (passed by ballot recently to become Version 3.0), vis:

  "A call originated on a Private UNI destined for an end system which only has a native (non-NSAP) E.164 address (i.e. a system directly attached to a public network supporting the native E.164 format) will code the Called Party number information element in the (NSAP) E.164 private ATM Address Format, with the RD, AREA, and ESI fields set to zero. The Called Party Subaddress information element is not used."

  Hence, in this case, ARP should return the E.164 address of the public ATM station in NSAP format. This is essentially implying an algorithmic resolution between the native E.164 and NSAP addresses of _directly_ attached public stations.
- Public network host to Public network host, no intervening private network: In this case, clearly the Q.93b requests would use native E.164 address formats.

- Public network host to Public network host, intervening private network: same as the case immediately above, since getting to and through the private network is a VC routing, not an addressing issue.

So several issues arise for ARP in supporting arbitrary connections between hosts on private and public network. One is how to distinguish between E.164 address and E.164 encoded NSAP modeled address. Another is what is the information to be supplied by ARP, e.g., in the public to private scenario should ARP return only the private NSAP modeled address or both an E.164 address, for a point of attachment between the public and private networks, along with the private NSAP modeled address.